

THE LARGE DEVIATION PRINCIPLE FOR THE ERDŐS-RÉNYI RANDOM GRAPH

SOURAV CHATTERJEE AND S. R. S. VARADHAN

ABSTRACT. What does an Erdős-Rényi graph look like when a rare event happens? This paper answers this question when p is fixed and n tends to infinity by establishing a large deviation principle under an appropriate topology. The formulation and proof of the main result uses the recent development of the theory of graph limits by Lovász and coauthors and Szemerédi's regularity lemma from graph theory. As a basic application of the general principle, we work out large deviations for the number of triangles in $G(n, p)$. Surprisingly, even this simple example yields an interesting double phase transition.

1. INTRODUCTION

1.1. The Erdős-Rényi graph. Let $G(n, p)$ be the random graph on n vertices where each edge is added independently with probability p . This model has been the subject of extensive investigations since the pioneering work of Erdős and Rényi [14], yielding a large body of literature (see [2, 18] for partial surveys).

This paper studies the following basic aspect of Erdős-Rényi graphs: What does the graph look like if one knows that some rare event has happened? One way to comprehensively answer this question is to formulate a large deviation principle for the Erdős-Rényi graph, in the same way as Sanov's theorem [27] gives a large deviation principle for an i.i.d. sample.

The setting of Sanov's theorem conforms naturally to the abstract theory of large deviations (see Chapter 6 in [11]) because i.i.d. samples can be thought of as random probability measures, allowing them to be viewed as random elements of a single topological space irrespective of the sample size. The first hurdle in formulating such a program for random graphs is in constructing a single abstract space in which all graphs can be embedded. Fortunately, this issue has been settled recently. In a sequence of papers [5, 6, 7, 15, 19, 20, 21, 22, 23, 24, 25] Laszlo Lovász with coauthors (listed here in order of frequency) V. T. Sós, B. Szegedy, C. Borgs, J. Chayes,

Key words and phrases. Random graph, Erdős-Rényi, graph limit, Szemerédi's lemma, large deviation principle.

Sourav Chatterjee's research was partially supported by NSF grants DMS-0707054 and DMS-1005312, and a Sloan Research Fellowship.

S. R. S. Varadhan's research was partially supported by NSF grants DMS-0904701 and OISE-0730136.

K. Vesztergombi, A. Schrijver and M. Freedman have developed a beautiful, unifying limit theory. (See also the related work of Diaconis and Janson [12] which traces this back to work of Aldous [1] and Hoover [17].) This sheds light on topics such as graph homomorphisms, Szemerédi’s regularity lemma, quasi-random graphs, graph testing and extremal graph theory, and has even found applications in statistics and related areas (see e.g. [10]). Their theory has been developed for dense graphs (number of edges comparable with the square of number of vertices) but parallel theories for sparse graphs are beginning to emerge [3].

1.2. Graph limits and graphons. The limit of a sequence of dense graphs can be defined as follows. We quote the definition verbatim from [22] (see also [6, 7, 12]). Let G_n be a sequence of simple graphs whose number of nodes tends to infinity. For every fixed simple graph H , let $|\text{hom}(H, G)|$ denote the number of homomorphisms of H into G (i.e. edge-preserving maps $V(H) \rightarrow V(G)$, where $V(H)$ and $V(G)$ are the vertex sets). This number is normalized to get the homomorphism density

$$(1) \quad t(H, G) := \frac{|\text{hom}(H, G)|}{|V(G)|^{|V(H)|}}.$$

This gives the probability that a random mapping $V(H) \rightarrow V(G)$ is a homomorphism.

Suppose that the graphs G_n become more and more similar in the sense that $t(H, G_n)$ tends to a limit $t(H)$ for every H . One way to define a limit of the sequence $\{G_n\}$ is to define an appropriate limit object from which the values $t(H)$ can be read off.

The main result of [22] (following the earlier equivalent work of Aldous [1] and Hoover [17]) is that indeed there is a natural “limit object” in the form of a function $f \in \mathcal{W}$, where \mathcal{W} is the space of all measurable functions from $[0, 1]^2$ into $[0, 1]$ that satisfy $f(x, y) = f(y, x)$ for all x, y .

Conversely, every such function arises as the limit of an appropriate graph sequence. This limit object determines all the limits of subgraph densities: if H is a simple graph with $V(H) = [k] = \{1, \dots, k\}$, then

$$(2) \quad t(H, f) = \int_{[0, 1]^k} \prod_{(i, j) \in E(H)} f(x_i, x_j) dx_1 \cdots dx_k.$$

Here $E(H)$ denotes the edge set of H . A sequence of graphs $\{G_n\}_{n \geq 1}$ is said to converge to f if for every finite simple graph H ,

$$(3) \quad \lim_{n \rightarrow \infty} t(H, G_n) = t(H, f).$$

Intuitively, the interval $[0, 1]$ represents a ‘continuum’ of vertices, and $f(x, y)$ denotes the probability of putting an edge between x and y . For example, for the Erdős-Rényi graph $G(n, p)$, if p is fixed and $n \rightarrow \infty$, then the limit graph is represented by the function that is identically equal to p on $[0, 1]^2$.

These limit objects, i.e. elements of \mathcal{W} , are called ‘graphons’ in [22, 6, 7]. A finite simple graph G on $\{1, \dots, n\}$ can also be represented as a graphon f^G in a natural way, by defining

$$(4) \quad f^G(x, y) = \begin{cases} 1 & \text{if } ([nx], [ny]) \text{ is an edge in } G, \\ 0 & \text{otherwise.} \end{cases}$$

Note that this allows *all* simple graphs, irrespective of the number of vertices, to be represented as elements of a single abstract space, namely \mathcal{W} .

1.3. The cut metric. With the above representation, it turns out that the notion of convergence in terms of subgraph densities outlined above can be captured by an explicit metric on \mathcal{W} , the so-called ‘cut distance’ (originally defined for finite graphs by Frieze and Kannan [16]). We start with the space \mathcal{W} of measurable functions $f(x, y)$ on $[0, 1]^2$ that satisfy $0 \leq f(x, y) \leq 1$ and $f(x, y) = f(y, x)$. We define the cut distance

$$(5) \quad d_{\square}(f, g) := \sup_{S, T \subseteq [0, 1]} \left| \int_{S \times T} [f(x, y) - g(x, y)] dx dy \right|.$$

We introduce in \mathcal{W} an equivalence relation. Let Σ be the space of measure preserving bijections $\sigma : [0, 1] \rightarrow [0, 1]$. Say that $f(x, y) \sim g(x, y)$ if $f(x, y) = g_{\sigma}(x, y) := g(\sigma x, \sigma y)$ for some $\sigma \in \Sigma$. Denote by \tilde{g} the closure in $(\mathcal{W}, d_{\square})$ of the orbit $\{g_{\sigma}\}$. The quotient space is denoted by $\tilde{\mathcal{W}}$ and τ denotes the natural map $g \rightarrow \tilde{g}$. Since d_{\square} is invariant under σ one can define on $\tilde{\mathcal{W}}$, the natural distance δ_{\square} by

$$\delta_{\square}(\tilde{f}, \tilde{g}) := \inf_{\sigma} d_{\square}(f, g_{\sigma}) = \inf_{\sigma} d_{\square}(f_{\sigma}, g) = \inf_{\sigma_1, \sigma_2} d_{\square}(f_{\sigma_1}, g_{\sigma_2})$$

making $(\tilde{\mathcal{W}}, \delta_{\square})$ into a metric space. To any finite graph G , we associate f^G as in (4) and its orbit $\tilde{G} = \tau f^G = \tilde{f}^G \in \tilde{\mathcal{W}}$. One of the key results of [6] is the following:

Theorem 1.1 (Theorem 3.8 in [6]). *A sequence of graphs $\{G_n\}_{n \geq 1}$ converges to a limit $f \in \mathcal{W}$ in the sense defined in (3) if and only if $\delta_{\square}(\tilde{G}_n, \tilde{f}) \rightarrow 0$ as $n \rightarrow \infty$.*

Szemerédi’s regularity lemma and the related deep results of Lovász and Szegedy will play a crucial role in this paper:

1.4. Szemerédi’s lemma. Let $G = (V, E)$ be a simple graph, and let X, Y be subsets of V . Then we denote by $e_G(X, Y)$ the number of X - Y edges of G (edges whose endpoints belong to $X \cap Y$ are counted twice), and call

$$\rho_G(X, Y) := \frac{e_G(X, Y)}{|X||Y|}$$

the *density* of the pair (X, Y) . Given some $\epsilon > 0$, we call a pair (A, B) of disjoint sets $A, B \subseteq V$ ϵ -regular if all $X \subseteq A$ and $Y \subseteq B$ with $|X| \geq \epsilon|A|$

and $|Y| \geq \epsilon|B|$ satisfy

$$|\rho_G(X, Y) - \rho_G(A, B)| \leq \epsilon.$$

A partition $\{V_0, \dots, V_K\}$ of V is called an ϵ -regular partition of G if it satisfies the following two conditions:

- (i) $|V_0| \leq \epsilon n$;
- (ii) $|V_1| = |V_2| = \dots = |V_K|$;
- (iii) all but at most ϵK^2 of the pairs (V_i, V_j) with $1 \leq i < j \leq K$ are ϵ -regular.

Szemerédi's regularity lemma goes as follows.

Theorem 1.2 (Szemerédi's lemma [28]). *Given $\epsilon > 0$ and an integer $m \geq 1$ there exists an integer $M = M(\epsilon, m)$ such that every graph of order at least M admits an ϵ -regular partition $\{V_0, \dots, V_K\}$ for some K in the range $m \leq K \leq M$.*

This result was proved by Szemerédi [28] in 1976 and has since found numerous applications in combinatorics, number theory and many other areas of discrete mathematics. The version presented above is from Diestel [13], Section 7.2. Lovász and Szegedy proved the following related result.

Theorem 1.3 (Theorem 5.1 in [23]). *The metric space $(\widetilde{\mathcal{W}}, \delta_\square)$ is compact.*

2. THE MAIN RESULT

2.1. The rate function. The main goal of this paper is to prove a large deviation principle for $G(n, p)$ when p is fixed and $n \rightarrow \infty$. The discussion in Section 1 gives a topological space (namely, $\widetilde{\mathcal{W}}$) that is suitable for this purpose. The next step is to define a rate function on this space. Let $I_p : [0, 1] \rightarrow \mathbb{R}$ be the function

$$\begin{aligned}
 (6) \quad I_p(u) &:= \frac{1}{2}u \log \frac{u}{p} + \frac{1}{2}(1-u) \log \frac{1-u}{1-p} \\
 &= \frac{1}{2} \sup_{a, b \in \mathbb{R}} \left[au + b(1-u) - \log (pe^a + (1-p)e^b) \right] \\
 &= \frac{1}{2} \sup_{a \in \mathbb{R}} \left[au - \log (pe^a + (1-p)) \right]
 \end{aligned}$$

The domain of the function I_p can be extended to \mathcal{W} as

$$\begin{aligned}
 I_p(h) &:= \int_0^1 \int_0^1 I_p(h(x, y)) dx dy \\
 (7) \quad &= \frac{1}{2} \int_0^1 \int_0^1 \left[h(x, y) \log \frac{h(x, y)}{p} + (1 - h(x, y)) \log \frac{1 - h(x, y)}{1 - p} \right] dx dy \\
 (8) \quad &= \frac{1}{2} \sup_{a(\cdot, \cdot)} \left[\int a(x, y) h(x, y) dx dy \right. \\
 &\quad \left. - \int \log(p e^{a(x, y)} + (1 - p)) dx dy \right].
 \end{aligned}$$

The following property of I_p is crucial.

Lemma 2.1. *The function I_p is well defined on $\widetilde{\mathcal{W}}$ and is lower semicontinuous under the cut metric δ_\square on $\widetilde{\mathcal{W}}$.*

Proof. The supremum in (8) can be taken over all bounded measurable functions a on $[0, 1]^2$. As the supremum of a family of affine linear functionals continuous in the weak topology, $I_p(h)$ is lower semi-continuous in the weak topology and therefore also in the topology of the metric d_\square . If $\sigma : [0, 1] \rightarrow [0, 1]$ is a measure preserving bijection then $I_p(h_\sigma) = I_p(h)$. By lower semi-continuity of $g \in \widetilde{h}$, $I_p(g) \leq I_p(h)$. But $g \in \widetilde{h}$ implies $h \in \widetilde{g}$ so that $I_p(h) \leq I_p(g)$. Hence, $I_p(g) = I_p(h)$ and $I_p(\cdot)$ is well defined and lower semi-continuous on $\widetilde{\mathcal{W}}$. \square

2.2. The Large Deviation Principle. The random graph $G(n, p)$ induces probability distributions $\mathbb{P}_{n,p}$ on the space \mathcal{W} through the map $G \rightarrow f^G$ and $\widetilde{\mathbb{P}}_{n,p}$ on $\widetilde{\mathcal{W}}$ through the map $G \rightarrow f^G \rightarrow \widetilde{f}^G = \tau f^G$. The space \mathcal{W} is compact in the weak topology and a large deviation principle for $\mathbb{P}_{n,p}$ on \mathcal{W} in the weak topology with the lower-semicontinuous rate function $I_p(h)$ given by (7) is elementary but is not of much use since quantities like ‘triangle counts’ are not stable in the weak topology. We will state it for the record and find a use for it later.

Theorem 2.2. *The sequence $\mathbb{P}_{n,p}$ on \mathcal{W} satisfies a large deviation principle in the weak topology. That is, for every weakly closed set $F \subset \mathcal{W}$*

$$\limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(F) \leq - \inf_{f \in F} I_p(f)$$

and for any open set U (again in the weak topology) in \mathcal{W}

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(U) \geq - \inf_{f \in U} I_p(f).$$

Proof. The weak topology is defined through an arbitrary but finite number of linear functionals. Therefore the large deviation principle can be reduced

to the large deviation behavior of a finite set of linear functionals $\{Z_\phi(f)\}$ given by

$$Z_\phi(f) = \langle \phi, f \rangle := \iint f(x, y) \phi(x, y) dx dy$$

under the measure $\mathbb{P}^{n,p}$. The limit

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{E}^{\mathbb{P}^{n,p}} \left[\exp \left(n^2 \iint f(x, y) \phi(x, y) dx dy \right) \right]$$

is easily calculated to yield

$$\frac{1}{2} \iint \log(p e^{2\phi(x,y)} + (1-p)) dx dy.$$

(Note that this is true only if ϕ is symmetric. However, since f is symmetric, it suffices to restrict attention to symmetric ϕ .) Therefore, an abstract Gärtner-Ellis Theorem (see e.g. Theorem 4.5.3 in [11]) gives the upper bound with rate function

$$I_p(f) = \sup_{\phi} \left[\langle \phi, f \rangle - \frac{1}{2} \iint \log(p e^{2\phi(x,y)} + (1-p)) dx dy \right].$$

Note that this is the rate function I_p defined in (7) and (8). The supremum is attained at the function

$$\phi_f(x, y) := \frac{1}{2} \log \frac{f(x, y)}{p} - \frac{1}{2} \log \frac{1 - f(x, y)}{1 - p}.$$

Note that for any $g \neq f$,

$$\begin{aligned} & (\langle \phi_f, f \rangle - I_p(f)) - (\langle \phi_f, g \rangle - I_p(g)) \\ &= \frac{1}{2} \iint \left(g(x, y) \log \frac{g(x, y)}{f(x, y)} + (1 - g(x, y)) \log \frac{1 - g(x, y)}{1 - f(x, y)} \right) dx dy > 0. \end{aligned}$$

This shows that every f is an exposed point of the lower semicontinuous rate function I_p , in the parlance of convex analysis. Therefore by the Gärtner-Ellis Theorem (see e.g. Theorem 4.5.20 in [11]) and the compactness of the weak topology, we get the lower bound. \square

The large deviation principle for $\tilde{\mathbb{P}}_{n,p}$ on $(\tilde{\mathcal{W}}, \delta_\square)$ is much more useful and is the main result of this article.

Theorem 2.3. *For each fixed $p \in (0, 1)$, the sequence $\tilde{\mathbb{P}}_{n,p}$ obeys a large deviation principle in the space $\tilde{\mathcal{W}}$ (equipped with the cut metric) with rate function I_p defined by (7). Explicitly, this means that for any closed set $\tilde{F} \subseteq \tilde{\mathcal{W}}$,*

$$(9) \quad \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \tilde{\mathbb{P}}_{n,p}(\tilde{F}) \leq - \inf_{\tilde{h} \in \tilde{F}} I_p(\tilde{h}).$$

and for any open set $\tilde{U} \subseteq \tilde{\mathcal{W}}$,

$$(10) \quad \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \tilde{\mathbb{P}}_{n,p}(\tilde{U}) \geq - \inf_{\tilde{h} \in \tilde{U}} I_p(\tilde{h}).$$

For the upper bound, because $(\widetilde{\mathcal{W}}, \delta_\square)$ is compact, it is sufficient to prove that for any $\tilde{h} \in \widetilde{\mathcal{W}}$,

$$\lim_{\eta \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \tilde{\mathbb{P}}_{n,p}(S_\square(\tilde{h}, \eta)) \leq -I_p(\tilde{h}).$$

For the lower bound we need to prove that if $\tilde{h} \in \widetilde{\mathcal{W}}$ and $\eta > 0$ is arbitrary

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \tilde{\mathbb{P}}_{n,p}(S_\square(\tilde{h}, \eta)) \geq -I_p(\tilde{h}),$$

where $S_\square(\tilde{h}, \eta) = \{\tilde{g} : \delta_\square(\tilde{g}, \tilde{h}) \leq \eta\}$.

2.3. Proof of the upper bound in Theorem 2.3. Let $B(\tilde{h}, \eta) \subset \mathcal{W}$ be defined as

$$B(\tilde{h}, \eta) = \tau^{-1} S_\square(\tilde{h}, \eta) \subset \mathcal{W}$$

i.e. the union of all the orbits from $S_\square(\tilde{h}, \eta) \subset \widetilde{\mathcal{W}}$. We need to show that

$$(11) \quad \lim_{\eta \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}[B(\tilde{h}, \eta)] \leq -I_p(\tilde{h})$$

Let the set of n vertices be partitioned into K subsets of size a with a remainder of size b , so that $n = Ka + b$. We assume that $b \leq \epsilon n$. We order the vertices so that $V_0 = \{1, 2, \dots, b\}$ and $V_i = \{b + (i-1)a + 1, \dots, b + ia\}$ for $i = 1, 2, \dots, K$. We map the vertices into subintervals of the unit interval, with the vertex r represented by the interval $[\frac{r-1}{n}, \frac{r}{n}]$. The sets V_i of vertices will then correspond to the intervals $E_0 = [0, \frac{b}{n}]$ for $i = 0$ and $E_i = [\frac{b+(i-1)a}{n}, \frac{b+ia}{n}]$ for $1 \leq i \leq K$. Let us denote by \mathcal{V}_K the subset of \mathcal{W} consisting of $g(x, y)$ that are of the form

$$g(x, y) = \sum_{i,j=1}^K p_{i,j} \mathbf{1}_{E_i}(x) \mathbf{1}_{E_j}(y)$$

where $\{p_{i,j}\}$, $1 \leq i, j \leq K$ is symmetric and satisfies $0 \leq p_{i,j} \leq 1$. For any pair m, M with $m < M$, we define $\mathcal{V}_{m,M} = \cup_{m \leq K \leq M} \mathcal{V}_K$. The following is a restatement of the Szemerédi regularity lemma.

Lemma 2.4. *Given any $\epsilon > 0$ and $m \geq 1$ such that $2m^{-1} < \epsilon$, there is $M = M(\epsilon, m)$ such that for any graph G , there exists a permutation π , i.e. a relabeling of the vertices of the graph, such that*

$$\inf_{g \in \mathcal{V}_{m,M}} d_\square(f^{\pi G}, g) \leq \epsilon.$$

Proof. Let ϵ' and m be given. According to Szemerédi's lemma, there is $M(\epsilon', m)$ such that, for some K in the range $m \leq K \leq M$, we can find a partition V_0, \dots, V_K which is ϵ' -regular. After a permutation we can assume that the ordering of the vertices coincides with the ordering of the partitions.

We define $p_{i,i} = p_{0,i} = p_{i,0} = 0$ and for $1 \leq i \neq j$, $p_{i,j} = \rho_G(V_i, V_j)$. This leads to

$$g(x, y) = \sum_{\substack{i,j=1 \\ i \neq j}}^K \rho_G(V_i, V_j) \mathbf{1}_{E_i}(x) \mathbf{1}_{E_j}(y)$$

to be compared with f^G when V_0, \dots, V_K is an ϵ' regular partition of G . Recall that

$$d_{\square}(f^G, g) = \sup_{S, T \subset [0,1]} \left| \iint_{S \times T} [f^G(x, y) - g(x, y)] dx dy \right|.$$

Since both f^G and g are constant on sets of the form $[\frac{i}{n}, \frac{i+1}{n}] \times [\frac{j}{n}, \frac{j+1}{n}]$ it is easy to see that S and T can be restricted to unions of intervals of the form $[\frac{i}{n}, \frac{i+1}{n}]$ i.e. subsets of $[0, 1]$ that represent subsets of vertices. These subsets will also be denoted by S and T . Now, given two such subsets S and T ,

$$\begin{aligned} & \iint_{S \times T} (f^G(x, y) - g(x, y)) dx dy \\ &= \sum_{0 \leq i, j \leq K} \iint_{(S \cap E_i) \times (T \cap E_j)} (f^G(x, y) - g(x, y)) dx dy. \end{aligned}$$

Let A_{ij} denote the (i, j) th term in the above sum. Since f^G and g both take values in $[0, 1]$, therefore for each i and each j ,

$$|A_{i0}| \leq |E_i| |E_0|, \quad |A_{0j}| \leq |E_0| |E_j|.$$

For the same reason, if $|S \cap E_i| < \epsilon' |E_i|$ or $|T \cap E_j| < \epsilon' |E_j|$, then

$$|A_{ij}| \leq \epsilon' |E_i| |E_j|.$$

If $1 \leq i \neq j$, $|S \cap E_i| \geq \epsilon' |E_i|$, $|T \cap E_j| \geq \epsilon' |E_j|$ and the pair (V_i, V_j) is ϵ' -regular, then

$$\begin{aligned} |A_{ij}| &= \left| \frac{1}{n^2} e_G(S \cap V_i, T \cap V_j) - \rho_G(V_i, V_j) |S \cap E_i| |T \cap E_j| \right| \\ &\leq |\rho_G(S \cap V_i, S \cap V_j) - \rho_G(V_i, V_j)| |S \cap E_i| |T \cap E_j| \\ &\leq \epsilon' |E_i| |E_j|. \end{aligned}$$

Finally, if either $1 \leq i = j$ or (V_i, V_j) is not ϵ' -regular (of which there are at most $K + 2\epsilon' K^2$ cases), then we have the trivial bound $|A_{ij}| \leq \frac{a^2}{n^2} \leq \frac{1}{K^2}$. A combination of the above estimates gives

$$\begin{aligned} \sum_{0 \leq i, j \leq K} |A_{ij}| &\leq 2|E_0| + \epsilon' + (K + 2\epsilon' K^2) \frac{1}{K^2} \\ &\leq (2\epsilon' + 2\epsilon' + \epsilon' + K^{-1}) \end{aligned}$$

Thus,

$$d_{\square}(f^G, g) \leq 5\epsilon' + K^{-1} \leq 5\epsilon' + m^{-1}$$

Since $m^{-1} < \epsilon/2$, we can choose ϵ' so that $5\epsilon' + m^{-1} < \epsilon$. \square

Lemma 2.5. *Let ϵ, m and M be as in Lemma 2.4.*

$$\mathbb{P}_{n,p}(B(\tilde{h}, \eta)) \leq n! \mathbb{P}_{n,p}(B(\tilde{h}, \eta) \cap B(\mathcal{V}_{m,M}, \epsilon))$$

where $B(\mathcal{V}_{m,M}, \epsilon) = \{g : \inf_{f \in \mathcal{V}_{m,M}} d_{\square}(g, f) \leq \epsilon\}$.

Proof. The orbit under the permutation group has at most $n!$ elements and they all have the same probability under $\mathbb{P}_{n,p}$. Moreover by the Lemma 2.4 every orbit meets $B(\mathcal{V}_{m,M}, \epsilon)$, and $B(\tilde{h}, \eta)$ is invariant under $\sigma \in \Sigma$ and therefore under π . Consequently

$$B(\tilde{h}, \eta) \subset \bigcup_{\pi} \pi^{-1}(B(\tilde{h}, \eta) \cap B(\mathcal{V}_{m,M}, \epsilon))$$

and the lemma follows. \square

Lemma 2.6. *There exists a function $\delta(\tilde{h}, \epsilon)$, depending only on \tilde{h} and ϵ , with $\delta(\tilde{h}, \epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$, such that for each arbitrary but fixed ϵ, m, M satisfying Lemma 2.4,*

$$\lim_{\eta \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(B(\tilde{h}, \eta) \cap B(\mathcal{V}_{m,M}, \epsilon)) \leq -I_p(\tilde{h}) + \delta(\tilde{h}, \epsilon)$$

Proof. Since $\mathcal{V}_{m,M}$ is a finite union $\bigcup_{m \leq K \leq M} \mathcal{V}_K$ it is sufficient to prove that for each K

$$\lim_{\eta \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(B(\tilde{h}, \eta) \cap B(\mathcal{V}_K, \epsilon)) \leq -I_p(\tilde{h}) + \delta(\tilde{h}, \epsilon)$$

and $\delta(\tilde{h}, \epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. For fixed K , \mathcal{V}_K consists of a compact set of functions in $L_1([0, 1]^2)$ and can be covered by a finite number of spheres of radius ϵ in L_1 and therefore in \mathcal{W} . It is therefore sufficient to show that for fixed K and $g \in \mathcal{V}_K$

$$\lim_{\eta \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(B(\tilde{h}, \eta) \cap B(g, 2\epsilon)) \leq -I_p(\tilde{h}) + \delta(\tilde{h}, \epsilon)$$

We can assume that $B(\tilde{h}, \eta) \cap B(g, 2\epsilon) \neq \emptyset$. Therefore $g \in B(\tilde{h}, \eta + 2\epsilon)$. Since $\eta \rightarrow 0$ we can assume $\eta < \epsilon$ so that $g \in B(\tilde{h}, 3\epsilon)$. By lower semi-continuity of $I_p(\cdot)$, $I_p(f) \geq I_p(\tilde{h}) - \delta(\tilde{h}, \epsilon)$ on $B(g, 2\epsilon) \subset B(\tilde{h}, 5\epsilon)$ and $\delta(\tilde{h}, \epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. We note that $B(g, 2\epsilon) \subset \mathcal{W}$ is weakly closed and therefore by the upper bound in Theorem 2.2,

$$\begin{aligned} \lim_{\eta \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(B(\tilde{h}, \eta) \cap B(g, 2\epsilon)) &\leq \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(B(g, 2\epsilon)) \\ &\leq - \inf_{f \in B(g, 2\epsilon)} I_p(f) \leq - \inf_{f \in B(\tilde{h}, 5\epsilon)} I_p(f) \leq -I_p(\tilde{h}) + \delta(\tilde{h}, \epsilon) \end{aligned}$$

where $\delta(\tilde{h}, \epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. \square

Lemma 2.5 and Lemma 2.6 yield (11), which proves the upper bound in Theorem 2.3.

2.4. Proof of the lower bound in Theorem 2.3. Let $h(x, y) \in \mathcal{W}$ be given. We define

$$p_{i,j}^{(n)} = n^2 \int \int_{[\frac{i-1}{n}, \frac{i}{n}] \times [\frac{j-1}{n}, \frac{j}{n}]} h(x, y) dx dy$$

and the corresponding function $h_n(x, y) \in \mathcal{W}$ by

$$h_n(x, y) = \sum_{i,j} p_{i,j}^{(n)} \mathbf{1}_{[\frac{i-1}{n}, \frac{i}{n}]}(x) \mathbf{1}_{[\frac{j-1}{n}, \frac{j}{n}]}(y).$$

Since $\|h_n - h\|_{L_1([0,1]^2)} \rightarrow 0$, it follows that $d_{\square}(h_n, h) \rightarrow 0$. It is therefore sufficient to prove that for any $\epsilon > 0$

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(d_{\square}(f^{G(n,p)}, h_n) \leq \epsilon) \geq -I_p(h).$$

We define an inhomogeneous random graph where the edge connecting the vertices i and j is present with probability $p_{i,j}^{(n)}$. Different edges are independent. If $\xi_{i,j} = 1$ when the edge connecting i, j is present and 0 otherwise then $\xi_{i,j}$ are independent Binomial random variables with $\mathbb{P}(\xi_{i,j} = 1) = p_{i,j}^{(n)}$. We denote by $\mathbb{P}_{n,h}$ the measure on \mathcal{W} induced by

$$f_n(x, y) := \sum_{\substack{i,j \\ i \neq j}} \xi_{i,j} \mathbf{1}_{[\frac{i-1}{n}, \frac{i}{n}]}(x) \mathbf{1}_{[\frac{j-1}{n}, \frac{j}{n}]}(y)$$

If A and B are subsets of $\{1, \dots, n\}$, it is straightforward to calculate

$$\begin{aligned} \psi_n(\lambda) &:= \frac{1}{n^2} \log \mathbb{E}^{\mathbb{P}_{n,h}} \left[\exp \left(\lambda \sum_{\substack{i \in A, j \in B \\ i \neq j}} (\xi_{i,j} - p_{i,j}^{(n)}) \right) \right] \\ &= \frac{1}{n^2} \left[\sum_{\substack{i,j \in A \cap B \\ i > j}} \log \mathbb{E}^{\mathbb{P}_{n,h}} \left[\exp(2\lambda(\xi_{i,j} - p_{i,j}^{(n)})) \right] \right. \\ &\quad \left. + \sum_{\substack{i \in A \cap B, j \in B \cap A^c \\ \text{or } i \in A \cap B^c, j \in A \cap B \\ \text{or } i \in A \cap B^c, j \in A^c \cap B}} \log \mathbb{E}^{\mathbb{P}_{n,h}} \left[\exp(\lambda(\xi_{i,j} - p_{i,j}^{(n)})) \right] \right]. \end{aligned}$$

Each term in the sum is easily estimated by $\frac{\lambda^2}{2}$, providing an estimate of the form

$$\mathbb{P}_{n,h} \left(\left| \iint_{\tilde{A} \times \tilde{B}} (f_n - h_n) dx dy \right| \geq \epsilon \right) \leq e^{-\frac{n^2 \epsilon^2}{2}},$$

where $\tilde{A} = \cup_{i \in A} [\frac{i-1}{n}, \frac{i}{n}]$ and \tilde{B} is defined similarly. Since the number of sets like $\tilde{A} \times \tilde{B}$ is only 2^{2n} it follows that

$$\mathbb{P}_{n,h}(d_{\square}(f_n, h_n) \geq \epsilon) \rightarrow 0$$

as $n \rightarrow \infty$. Now the lower bound is easily established by a simple tilting argument. Denoting by $B_{\epsilon,n}$ the set $\{f : d_{\square}(f, h_n) \leq \epsilon\}$

$$\begin{aligned} \mathbb{P}_{n,p}(B_{\epsilon,n}) &= \int_{B_{\epsilon,n}} d\mathbb{P}_{n,p} = \int_{B_{\epsilon,n}} e^{-\log \frac{d\mathbb{P}_{n,h}}{d\mathbb{P}_{n,p}}} d\mathbb{P}_{n,h} \\ &= \mathbb{P}_{n,h}(B_{\epsilon,n}) \frac{1}{\mathbb{P}_{n,h}(B_{\epsilon,n})} \int_{B_{\epsilon,n}} e^{-\log \frac{d\mathbb{P}_{n,h}}{d\mathbb{P}_{n,p}}} d\mathbb{P}_{n,h}. \end{aligned}$$

By Jensen's inequality

$$\log \mathbb{P}_{n,p}(B_{\epsilon,n}) \geq \log \mathbb{P}_{n,h}(B_{\epsilon,n}) - \frac{1}{\mathbb{P}_{n,h}(B_{\epsilon,n})} \int_{B_{\epsilon,n}} \log \frac{d\mathbb{P}_{n,h}}{d\mathbb{P}_{n,p}} d\mathbb{P}_{n,h}.$$

Since $\mathbb{P}_{n,h}(B_{\epsilon,n}) \rightarrow 1$, it is easy to see that

$$\liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}_{n,p}(B_{\epsilon,n}) \geq - \lim_{n \rightarrow \infty} \frac{1}{n^2} \int \log \frac{d\mathbb{P}_{n,h}}{d\mathbb{P}_{n,p}} d\mathbb{P}_{n,h}$$

The entropy cost of tilting (i.e. the integral in the preceding display) is

$$\frac{1}{n^2} \sum_{i>j} \left(p_{i,j}^{(n)} \log \frac{p_{i,j}^{(n)}}{p} + (1 - p_{i,j}^{(n)}) \log \frac{1 - p_{i,j}^{(n)}}{1 - p} \right)$$

which converges to $I_p(h)$ as $n \rightarrow \infty$. This proves the lower bound.

3. CONDITIONAL DISTRIBUTIONS

Theorem 2.3 gives estimates of the probabilities of rare events for $G(n, p)$. However, it does not quite answer the following question: given that some particular rare event has occurred, what does the graph look like? Naturally, one might expect that if $G(n, p) \in \tilde{F}$ for some closed set $\tilde{F} \subseteq \tilde{\mathcal{W}}$ satisfying

$$(12) \quad \inf_{\tilde{h} \in \tilde{F}^o} I_p(\tilde{h}) = \inf_{\tilde{h} \in \tilde{F}} I_p(\tilde{h}) > 0,$$

then $G(n, p)$ should resemble one of the minimizers of I_p in \tilde{F} . (Here \tilde{F}^o denotes the interior of \tilde{F} , as usual.) In other words, given that $G(n, p) \in \tilde{F}$, one might expect that $\delta_{\square}(G(n, p), \tilde{F}^*) \approx 0$, where \tilde{F}^* is the set of minimizers of I_p in \tilde{F} and

$$\delta_{\square}(G(n, p), \tilde{F}^*) := \inf_{\tilde{h} \in \tilde{F}^*} \delta_{\square}(G(n, p), \tilde{h}).$$

However, it is not obvious that a minimizer must exist in \tilde{F} . Here is where the compactness of $\tilde{\mathcal{W}}$ comes to the rescue yet one more time: since the function I_p is lower semicontinuous on \tilde{F} and \tilde{F} is closed, therefore a minimizer must necessarily exist. The following theorem formalizes this argument.

Theorem 3.1. *Take any $p \in (0, 1)$. Let \tilde{F} be a closed subset of $\tilde{\mathcal{W}}$ satisfying (12). Let \tilde{F}^* be the subset of \tilde{F} where I_p is minimized. Then \tilde{F}^* is non-empty and compact, and for each n , and each $\epsilon > 0$,*

$$\mathbb{P}(\delta_{\square}(G(n, p), \tilde{F}^*) \geq \epsilon \mid G(n, p) \in \tilde{F}) \leq e^{-C(\epsilon, \tilde{F})n^2}$$

where $C(\epsilon, \tilde{F})$ is a positive constant depending only on ϵ and \tilde{F} . In particular, if \tilde{F}^* contains only one element \tilde{h}^* , then the conditional distribution of $G(n, p)$ given $G(n, p) \in \tilde{F}$ converges to the point mass at \tilde{h}^* as $n \rightarrow \infty$.

Proof. Since $\tilde{\mathcal{W}}$ is compact and \tilde{F} is a closed subset, therefore \tilde{F} is also compact. Since I_p is a lower semicontinuous function on \tilde{F} (Lemma 2.1) and \tilde{F} is compact, it must attain its minimum on \tilde{F} . Thus, \tilde{F}^* is non-empty. By the lower semicontinuity of I_p , \tilde{F}^* is closed (and hence compact). Fix $\epsilon > 0$ and let

$$\tilde{F}_{\epsilon} := \{\tilde{h} \in \tilde{F} : \delta_{\square}(\tilde{h}, \tilde{F}^*) \geq \epsilon\}.$$

Then \tilde{F}_{ϵ} is again a closed subset. Observe that

$$\mathbb{P}(\delta_{\square}(G(n, p), \tilde{F}^*) \geq \epsilon \mid G(n, p) \in \tilde{F}) = \frac{\mathbb{P}(G(n, p) \in \tilde{F}_{\epsilon})}{\mathbb{P}(G(n, p) \in \tilde{F})}.$$

Thus, with

$$I_1 := \inf_{\tilde{h} \in \tilde{F}} I_p(\tilde{h}), \quad I_2 := \inf_{\tilde{h} \in \tilde{F}_{\epsilon}} I_p(\tilde{h}),$$

Theorem 2.3 and condition (12) give

$$\limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(\delta_{\square}(G(n, p), \tilde{F}^*) \geq \epsilon \mid G(n, p) \in \tilde{F}) \leq I_1 - I_2.$$

The proof will be complete if it is shown that $I_1 < I_2$.

Now clearly, $I_1 \leq I_2$. If $I_1 = I_2$, the compactness of \tilde{F}_{ϵ} implies that there exists $\tilde{h} \in \tilde{F}_{\epsilon}$ satisfying $I_p(\tilde{h}) = I_2$. However, this means that $\tilde{h} \in \tilde{F}^*$ and hence $\tilde{F}_{\epsilon} \cap \tilde{F}^* \neq \emptyset$, which is impossible. \square

4. APPLICATION TO TRIANGLE COUNTS

4.1. Brief history of the problem. Let $T_{n,p}$ be the number of triangles in $G(n, p)$. The primary objective of this section is to compute the large deviation rate function for the upper tail of $T_{n,p}$ when p remains fixed and $n \rightarrow \infty$. In other words, given $p \in [0, 1]$ and $\epsilon > 0$, we wish to evaluate the limit

$$(13) \quad \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(T_{n,p} \geq (1 + \epsilon)\mathbb{E}(T_{n,p}))$$

as a function of p and ϵ .

The problem of estimating tail probabilities like $\mathbb{P}(T_{n,p} \geq (1 + \epsilon)\mathbb{E}(T_{n,p}))$ has been studied extensively in the random graphs literature, particularly in the case when p is allowed to tend to zero as $n \rightarrow \infty$. Computing upper and lower bounds on such tail probabilities that are sharp up to constants

in the exponent was a prominent open problem in this area until it was solved recently in [8]. Let us refer to the paper [8] for a survey of the aforementioned literature.

When p is fixed, computing sharp upper and lower bounds is relatively easy. The difficult problem in this case is the exact evaluation of the limit (13). The first progress in this direction was made in [9] where it was shown that, given $p \in (0, 1)$, there exist $p^3/6 < t' \leq t'' < 1/6$ such that for all $t \in (p^3/6, t') \cup (t'', 1/6)$,

$$(14) \quad \lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(T_{n,p} \geq tn^3) = -I_p((6t)^{1/3}),$$

when I_p is the entropy function defined in (6). Explicit formulas for p' and p'' are also given in [9]. Unfortunately, the result does not cover all values of (p, t) ; and neither is the above formula true for all (p, t) , as we shall see below.

There is a related unpublished manuscript by Bolthausen, Comets and Dembo [4] on large deviations for subgraph counts. As of now, to the best of our knowledge, the authors of [4] have only looked at subgraphs that do not complete loops, like 2-stars. Besides [9] and [4], we know of no other papers that attack the exact evaluation of (13) (or equivalently, (14)).

4.2. Exact large deviations for the upper tail. In this subsection, the limit (14) is evaluated for all values of p and t . It comes as the solution of the following variational problem. Let \mathcal{W} , $\widetilde{\mathcal{W}}$ and δ_\square be defined as in Section 1. For each $f \in \mathcal{W}$, let

$$T(f) := \frac{1}{6} \int_0^1 \int_0^1 \int_0^1 f(x, y) f(y, z) f(z, x) dx dy dz$$

and let $I_p(f)$ be defined as in (7). Note that T can be defined on $\widetilde{\mathcal{W}}$ simply by letting $T(\tilde{f}) := T(f)$, because T is a continuous map on \mathcal{W} under the d_\square pseudometric (Theorem 3.7 in [6]).

For each $p \in [0, 1]$ and $t \in [0, 1/6)$, let

$$(15) \quad \phi(p, t) := \inf\{I_p(f) : f \in \mathcal{W}, T(f) \geq t\}.$$

For $t \geq 1/6$, let $\phi(p, t) = \infty$. The following result gives the large deviation rate function for the upper tail of $T_{n,p}$. (Note that this is just an illustrative example. Theorem 2.3 can be used to derive large deviations for any subgraph count, or even joint large deviations for the counts of more than one subgraph.)

Theorem 4.1. *Let $G(n, p)$ be the Erdős-Rényi random graph on n vertices with edge probability p . Let $T_{n,p}$ denote the number of triangles in $G(n, p)$. Let ϕ be defined as above. Then for each $p \in (0, 1)$ and each $t \geq 0$,*

$$\lim_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(T_{n,p} \geq tn^3) = -\phi(p, t).$$

Next, take any $p \in (0, 1)$ and $t \in (p^3/6, 1/6)$. Let $F_{p,t}^*$ be the set of minimizers for the variational problem (15) and $\tilde{F}_{p,t}^*$ be its image in $\widetilde{\mathcal{W}}$. Then $\tilde{F}_{p,t}^*$ is a non-empty compact subset of $\widetilde{\mathcal{W}}$. Moreover, for each $\epsilon > 0$ there exists a positive constant $C(\epsilon, p, t)$ depending only on ϵ , p and t such that for any n ,

$$\mathbb{P}(\delta_{\square}(G(n, p), \tilde{F}_{p,t}^*) \geq \epsilon \mid T_{n,p} \geq tn^3) \leq e^{-C(\epsilon, p, t)n^2}.$$

Proof. Let $F := \{f \in \mathcal{W} : T(f) \geq t\}$. By Theorem 3.7 in [6], F is a closed subset of \mathcal{W} . Therefore by Theorem 2.3,

$$\begin{aligned} \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(T_{n,p} \geq tn^3) &= \limsup_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(G(n, p) \in F) \\ &\leq -\inf_{h \in F} I_p(h) = -\phi(p, t). \end{aligned}$$

Next, let $U := \{f \in \mathcal{W} : T(f) > t\}$. Again by Theorem 3.7 of [6], U is an open set. Therefore by Theorem 2.3, for each $\epsilon > 0$,

$$\begin{aligned} \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(T_{n,p} \geq tn^3) &\geq \liminf_{n \rightarrow \infty} \frac{1}{n^2} \log \mathbb{P}(G(n, p) \in U) \\ &\geq -\inf_{h \in U} I_p(h) \geq -\phi(p, t + \epsilon). \end{aligned}$$

In Proposition 4.2 below, it is proved that ϕ is a continuous function of t for every fixed p . This completes the proof of the first assertion of the theorem. The second assertion is merely a corollary of Theorem 3.1. The condition (12) required for Theorem 3.1 can be easily shown to follow from the continuity of ϕ in t , because any \tilde{f} with $T(\tilde{f}) > t$ lies in the interior of the set $\{\tilde{h} : T(\tilde{h}) \geq t\}$. \square

4.3. Properties of the rate function. Given Theorem 2.3, there is a natural desire to understand the rate function ϕ . The following proposition summarizes some basic properties of ϕ . The first property is required in the proof of Theorem 4.1 above.

Proposition 4.2. *For each fixed $p \in (0, 1)$, the following hold:*

- (i) *The function $\phi(p, t)$ is continuous in t in the interval $[0, 1/6]$.*
- (ii) *As a function of t , $\phi(p, t) = 0$ in the interval $[0, p^3/6]$ and strictly increasing in $(p^3/6, 1/6)$. Moreover, for $p^3/6 < t < s < 1/6$,*

$$\phi(p, t) < (t/s)^{1/3} \phi(p, s).$$

- (iii) *For $t \in (p^3/6, 1/6)$, $\phi(p, t)$ can be alternately represented as*

$$\phi(p, t) := \inf\{I_p(f) : f \in \mathcal{W}, T(f) = t\}.$$

Moreover, if $\{f_n\}_{n \geq 1}$ is a sequence in \mathcal{W} such that $T(f_n) \geq t$ for all n and $I_p(f_n) \rightarrow \phi(p, t)$, then $T(f_n) \rightarrow t$. In particular, the elements of $\tilde{F}_{p,t}^$ all satisfy $T(f) = t$.*

Proof. Let us start by proving that ϕ is continuous in t . For each $f \in \mathcal{W}$ and $\delta \in [0, 1]$, let

$$f^\delta := f + \delta(1 - f).$$

By the inequality

$$\begin{aligned} & (a + \delta(1 - a))(b + \delta(1 - b))(c + \delta(1 - c)) \\ &= abc + \delta((1 - a)bc + (1 - b)ac + (1 - c)ab) \\ & \quad + \delta^2((1 - a)(1 - b)c + (1 - a)(1 - c)b + (1 - b)(1 - c)a) \\ & \quad + \delta^3(1 - a)(1 - b)(1 - c) \\ &\geq abc + \delta^3((1 - a)bc + (1 - b)ac + (1 - c)ab) \\ & \quad + \delta^3((1 - a)(1 - b)c + (1 - a)(1 - c)b + (1 - b)(1 - c)a) \\ & \quad + \delta^3(1 - a)(1 - b)(1 - c) \\ &= abc + \delta^3(1 - abc), \end{aligned}$$

we see that

$$(16) \quad T(f^\delta) \geq T(f)(1 - \delta^3) + \frac{\delta^3}{6}.$$

Take any $t \in [0, 1/6]$ and any f such that $T(f) \geq t$. Suppose $t_n \downarrow t$. Let δ_n be the smallest number in $[0, 1]$ such that $T(f^{\delta_n}) \geq t_n$. By (16) it follows that δ_n exists and $\lim_{n \rightarrow \infty} \delta_n = 0$. Therefore by the dominated convergence theorem, $\lim_{n \rightarrow \infty} I_p(f^{\delta_n}) = I_p(f)$. Thus,

$$\lim_{n \rightarrow \infty} \phi(p, t_n) \leq \lim_{n \rightarrow \infty} I_p(f^{\delta_n}) = I_p(f).$$

Since this is true for every f such that $T(f) \geq t$ and ϕ is a non-decreasing function in t , this proves the right continuity of ϕ .

Next, take a sequence $t_n \uparrow t$. Let f_n be a sequence of functions such that $T(f_n) \geq t_n$ and

$$\lim_{n \rightarrow \infty} I_p(f_n) = \lim_{n \rightarrow \infty} \phi(p, t_n).$$

For each n , let δ_n be the smallest number in $(0, 1)$ such that $T(f_n^{\delta_n}) \geq t$. By (16), δ_n exists and $\lim_{n \rightarrow \infty} \delta_n = 0$. Now, the function I_p on $[0, 1]$ (defined in (6)) is uniformly continuous on $[0, 1]$. As a consequence,

$$\lim_{\delta \rightarrow 0} \sup_{f \in \mathcal{W}} |I_p(f^\delta) - I_p(f)| = 0.$$

In particular,

$$\lim_{n \rightarrow \infty} I_p(f_n^{\delta_n}) = \lim_{n \rightarrow \infty} \phi(p, t_n).$$

But $\phi(p, t) \leq I_p(f_n^{\delta_n})$ for each n . By the monotonicity of ϕ , this proves left continuity.

Next, note that since $\mathbb{P}(T(G_{n,p}) \geq tn^3) \rightarrow 1$ for any $t < p^3/6$, Theorem 4.1 and the continuity of ϕ imply that $\phi(p, t) = 0$ for $t \leq p^3/6$. Let us now show that $\phi(p, t)$ is strictly increasing in t when $t \in (p^3/6, 1/6)$.

Fix $p^3/6 \leq t < s < 1/6$. Fix $\epsilon > 0$. Take any $f \in \mathcal{W}$ such that $T(f) \geq s$. For each $\delta \in (0, 1)$, let

$$(17) \quad f_\delta(x, y) := (1 - \delta)f(x, y) + \delta p.$$

By the inequality

$$((1 - \delta)a + \delta p)((1 - \delta)b + \delta p)((1 - \delta)c + \delta p) \geq (1 - \delta)^3 abc$$

we have

$$T(f_\delta) \geq (1 - \delta)^3 T(f) \geq (1 - \delta)^3 s.$$

Thus, if we take δ such that $(1 - \delta)^3 = t/s$, then $T(f_\delta) \geq t$. By the convexity of the function I_p on $[0, 1]$ defined in (6), we see that for any $x \geq p$,

$$I_p((1 - \delta)x + \delta p) \leq (1 - \delta)I_p(x) + \delta I_p(p) = (1 - \delta)I_p(x),$$

and therefore,

$$(18) \quad I_p(f_\delta) \leq (1 - \delta)I_p(f).$$

Thus, $\phi(p, t) \leq I_p(f_\delta) \leq (t/s)^{1/3} I_p(f)$. Since this holds for any f with $T(f) \geq s$, it shows that

$$(19) \quad \phi(p, t) \leq (t/s)^{1/3} \phi(p, s).$$

To show ϕ is a strictly increasing function of t in the interval $(p^3/6, 1/6)$, it therefore suffices to prove that $\phi(p, t) > 0$ for t in this interval. This follows easily, since the strict convexity of I_p on $[0, 1]$ and equation (18) show that equality in (19) can hold only if $f \equiv p$ almost everywhere for some f such that $T(f) \geq s$, which is impossible since $s > p^3/6$.

Next, fix $t \in (p^3/6, 1/6)$ and take any sequence $\{f^{(n)}\}$ in \mathcal{W} such that $T(f^{(n)}) \geq t$ for all n and $I_p(f^{(n)}) \rightarrow \phi(p, t)$. Recall the subscript notation introduced in (17) above. By the continuous mapping theorem, there exist $\delta_n \in [0, 1]$ such that for each n ,

$$T(f_{\delta_n}^{(n)}) = t.$$

Therefore by (18),

$$\phi(p, t) \leq I_p(f_{\delta_n}^{(n)}) \leq (1 - \delta_n)I_p(f^{(n)}),$$

which proves that $\delta_n \rightarrow 0$. This proves that

$$\phi(p, t) = \inf\{I_p(f) : f \in \mathcal{W}, T(f) = t\}.$$

Since $\delta_n \rightarrow 0$, this also proves that $T(f^{(n)}) \rightarrow t$. This completes the proof of Proposition 4.2. \square

4.4. The ‘Replica Symmetric’ phase. Note that there are two “extreme” functions that satisfy $T(f) = t$. First, there is the constant function

$$(20) \quad c_t(x, y) \equiv (6t)^{1/3}.$$

On the opposite extreme, there is the function χ_t , defined as

$$(21) \quad \chi_t(x, y) := \begin{cases} 1 & \text{if } \max\{x, y\} \leq (6t)^{1/3}, \\ 0 & \text{otherwise.} \end{cases}$$

In a limiting sense, c_t represents an Erdős-Rényi random graph with edge probability $(6t)^{1/3}$, while χ_t represents the union of a clique of size $n(6t)^{1/3}$ and a set of isolated vertices of size $n(1 - (6t)^{1/3})$.

It is simple to see that (14) holds if and only if the infimum in (15) is attained at the constant function c_t . The following theorem gives a sufficient condition for this to happen. This extends the main result of [9]. The methods of [9] are closely related to methods from statistical physics; drawing inspiration from this connection, one may call the region where c_t solves the variational problem (15) as the ‘replica symmetric phase’ of the problem.

Theorem 4.3. *For each $0 < p < 1$ and $t \in (p^3/6, 1/6)$, let $h_p(t) := I_p((6t)^{1/3})$. For $t \in [0, p^3/6]$, let $h_p(t) = 0$. Let \hat{h}_p be the convex minorant of h_p (i.e. the maximum convex function lying below h_p). If t is a point in $(p^3/6, 1/6)$ where $h_p(t) = \hat{h}_p(t)$, then the variational problem (15) for this pair (p, t) is uniquely solved by the constant function $c_t \equiv (6t)^{1/3}$. Consequently, $\phi(p, t) = h_p(t)$. Moreover, for such (p, t) , for each $\epsilon > 0$*

$$\lim_{n \rightarrow \infty} \mathbb{P}(\delta_{\square}(G(n, p), c_t) \geq \epsilon \mid T_{n,p} \geq tn^3) = 0.$$

Since c_t is the limit of $G(n, (6t)^{1/3})$, this means that for such (p, t) the conditional distribution of $G(n, p)$ given $T_{n,p} \geq tn^3$ is indistinguishable from the law of $G(n, (6t)^{1/3})$ in the large n limit.

Proof. Since h_p is an increasing function, it is easy to see that $h_p(t) = \hat{h}_p(t)$ for $t \leq p^3/6$. (Incidentally, this also shows that \hat{h}_p is an increasing function in $[0, 1/6)$, and strictly increasing in $(p^3/6, 1/6)$.)

Suppose t is a point in $(p^3/6, 1/6)$ such that $h_p(t) = \hat{h}_p(t)$. We claim that there exists $\beta > 0$ such that

$$(22) \quad t = \operatorname{argmax}_{x \in [0, 1/6]} (\beta x - h_p(x)).$$

To see this, observe that since \hat{h}_p is convex and strictly increasing in the interval $(p^3/6, 1/6)$, there exists $\beta > 0$ and $c \in \mathbb{R}$ such that the line $y = \beta x + c$ lies below the curve $y = \hat{h}_p(x)$ and touches it at $x = t$. But we also know that h_p lies above \hat{h}_p and the two curves touch at t . Thus,

$$t = \operatorname{argmax}_x (\beta x + c - h_p(x)) = \operatorname{argmax}_x (\beta x - h_p(x)).$$

This proves the claim.

Now take any $f \in \mathcal{W}$ such that $T(f) \geq t$. Let c_t be the function that is identically equal to $(6t)^{1/3}$. Then $T(c_t) = t$, and by Hölder's inequality and (22),

$$\begin{aligned} \beta t - I_p(c_t) &= \iint \left(\frac{\beta c_t(x, y)^3}{6} - I_p(c_t(x, y)) \right) dx dy \\ &\geq \iint \left(\frac{\beta f(x, y)^3}{6} - I_p(f(x, y)) \right) dx dy \\ &\geq \frac{\beta}{6} \iiint f(x, y) f(y, z) f(z, x) dx dy dz - \iint I_p(f(x, y)) dx dy \\ &\geq \beta t - I_p(f). \end{aligned}$$

Thus, c_t minimizes $I_p(f)$ among all f such that $T(f) \geq t$. This shows that

$$\phi(p, t) = I_p(c_t) = h_p(t).$$

The uniqueness of the optimizer follows from the Hölder step in the above deduction. Finally, the claim about the conditional distribution follows from Theorem 3.1 and the uniqueness of the minimizer. \square

It is easy to show that for any $p > 0$, $h_p(t) = \hat{h}_p(t)$ for all $t \in (p^3/6, t') \cup (t'', 1/6)$ where t' and t'' depend on p . Similarly, given any $t \in (0, 1/6)$, there exists $p' < (6t)^{1/3}$ depending on t such that for all $p \in (p', (6t)^{1/3})$, $h_p(t) = \hat{h}_p(t)$. Thus, there is a nontrivial set of (p, t) where c_t solves the variational problem and consequently $\phi(p, t) = h_p(t)$. As mentioned before, this recovers the main result of [9]. The conclusion about the conditional distribution is a new result.

4.5. Replica Symmetry Breaking. Given Theorem 4.3, it is quite interesting to note that the variational problem (15) is not solved by constant functions everywhere. From the physical point of view espoused in [9], however, this is not surprising; it is simply the effect of replica symmetry breaking down in the ‘low temperature regime’.

The phase transition is very easy to establish using Theorem 4.1, by comparing the performances of c_t and χ_t defined in (20) and (21). A simple computation shows that for any $t \in (0, 1/6)$,

$$(23) \quad \lim_{p \rightarrow 0} \frac{I_p(c_t)}{\log(1/p)} = \frac{(6t)^{1/3}}{2} > \frac{(6t)^{2/3}}{2} = \lim_{p \rightarrow 0} \frac{I_p(\chi_t)}{\log(1/p)}.$$

Combining the above observation with Theorem 3.1, it follows easily that there are values of (p, t) such that given $T_{n,p} \geq tn^3$, the graph $G(n, p)$ must look different than an Erdős-Rényi graph. Again, compactness is crucial. This is formalized by the following theorem.

Theorem 4.4. *Let $\tilde{\mathcal{C}}$ denote the set of constant functions in $\tilde{\mathcal{W}}$. For each $t \in (0, 1/6)$, there exists $p' > 0$ such that for all $p < p'$, the variational problem (15) is not solved by the constant function $c_t \equiv (6t)^{1/3}$. Moreover,*

for such (p, t) , $\delta_{\square}(\tilde{F}_{p,t}^*, \tilde{C}) > 0$, where $\tilde{F}_{p,t}^*$ is the set of minimizers defined in Theorem 4.1. Consequently, there exists $\epsilon > 0$ such that

$$\lim_{n \rightarrow \infty} \mathbb{P}(\delta_{\square}(G(n, p), \tilde{C}) > \epsilon \mid T_{n,p} \geq tn^3) = 1.$$

Proof. By (23) we see that for each t , there exists $p' > 0$ such that c_t is not a minimizer for the problem (15) if $p < p'$. Take any such (p, t) . By Theorem 4.1, $\tilde{F}_{p,t}^*$ is non-empty and compact. By part (iii) of Proposition 4.2 and the non-optimality of c_t , it follows that \tilde{C} and $\tilde{F}_{p,t}^*$ must be disjoint. But \tilde{C} and $\tilde{F}_{p,t}^*$ are both compact subsets of $\tilde{\mathcal{W}}$. Therefore, $\delta_{\square}(\tilde{C}, \tilde{F}_{p,t}^*) > 0$. The last claim follows by Theorem 4.1. \square

4.6. The double phase transition. A combination of Theorem 4.4 and Theorem 4.3 shows that for all small enough p , the variational problem (15) has a ‘double phase transition’. (Actually, there may be more than two phase transitions, but we show that there is at least two.)

Indeed, for all small enough p , Theorem 4.3 (or the results of [9]) show that there exists $p^3/6 < t' \leq t'' < 1/6$ such that $\phi(p, t) = I_p((6t)^{1/3})$ for all $t \in (p^3/6, t') \cup (t'', 1/6)$. On the other hand by Theorem 4.4, it follows that for all small enough p , the variational problem (15) is not solved by a constant function at the point $(p, 1/2)$. Combining these two observations gives the following theorem.

Theorem 4.5. *There exists $p_0 > 0$ such that if $p \leq p_0$, then there exists $p^3/6 < t' < t'' < 1/6$ such that the variational problem (15) is solved by the constant function $c_t \equiv (6t)^{1/3}$ when $t \in (p^3/6, t') \cup (t'', 1/6)$, but there is a non-empty subset of (t', t'') where all optimizers are non-constant.*

Of course, as shown by Theorems 4.3 and 4.4, the significance of optimizers being constant or non-constant is in whether the conditional behavior of $G(n, p)$ given $T_{n,p} \geq tn^3$ is close to that of an Erdős-Rényi graph or not.

4.7. The small p limit. The last theorem of this paper describes the nature of $\phi(p, t)$ and $\tilde{F}_{p,t}^*$ when t is fixed and p is very small, tending to zero. The essence of the result, perhaps not surprisingly, is that when t is fixed and $p \rightarrow 0$, then conditionally on the event $\{T_{n,p} \geq tn^3\}$ the graph $G(n, p)$ must look like a clique.

Theorem 4.6. *For each $t \in (0, 1/6)$,*

$$\lim_{p \rightarrow 0} \frac{\phi(p, t)}{\log(1/p)} = \frac{(6t)^{2/3}}{2}.$$

Moreover, if χ_t is the function defined in (21) and $\tilde{F}_{p,t}^$ is defined as in Theorem 4.1, then the set $\tilde{F}_{p,t}^*$ converges to the point $\tilde{\chi}_t$ as $p \rightarrow 0$, in the sense that*

$$\lim_{p \rightarrow 0} \sup_{\tilde{f} \in \tilde{F}_{p,t}^*} \delta_{\square}(\tilde{f}, \tilde{\chi}_t) = 0.$$

Consequently, for each $\epsilon > 0$,

$$\lim_{p \rightarrow 0} \lim_{n \rightarrow \infty} \mathbb{P}(\delta_{\square}(G(n, p), \tilde{\chi}_t) \geq \epsilon \mid T_{n,p} \geq tn^3) = 0.$$

Proof. In this proof, C will denote any constant that does not depend on anything else. The value of C may change from line to line. All integrals will be over the interval $[0, 1]$.

Fix $t \in (0, 1/6)$. For each $p < (6t)^{1/3}$, choose a function $f_p \in F_{p,t}^*$. From the definition of I_p , observe that if $p \leq 1/2$,

$$(24) \quad \left| I_p(f_p) - \frac{1}{2} \iint f_p(x, y) \log \frac{1}{p} dx dy \right| \leq C.$$

On the other hand, by the definition of $\tilde{F}_{p,t}^*$,

$$I_p(f_p) \leq I_p(\chi_t) \leq \frac{(6t)^{2/3}}{2} \log \frac{1}{p} + C.$$

Combining the last two inequalities gives

$$(25) \quad \iint f_p(x, y) dx dy \leq (6t)^{2/3} + \frac{C}{\log(1/p)}.$$

Next, let

$$h_p(x, y) := \int f_p(x, z) f_p(z, y) dz.$$

Then by two applications of the Cauchy-Schwarz inequality, the inequality (25), and the fact that $f_p(x, y) \in [0, 1]$, we get the following important sequence of inequalities.

$$(26) \quad \left\{ \begin{aligned} (6t)^2 &= (6T(f_p))^2 = \left(\iint h_p(x, y) f_p(x, y) dx dy \right)^2 \\ &\leq \iint h_p^2(x, y) dx dy \iint f_p^2(x, y) dx dy \\ &\leq \int \left(\int f_p^2(x, z) dz \int f_p^2(y, z) dz \right) dx dy \iint f_p^2(x, y) dx dy \\ &= \left(\iint f_p^2(x, y) dx dy \right)^3 \\ &\leq \left(\iint f_p(x, y) dx dy \right)^3 \leq \left((6t)^{2/3} + \frac{C}{\log(1/p)} \right)^3. \end{aligned} \right.$$

A direct consequence of (26), combined with (24), is that

$$\lim_{p \rightarrow 0} \frac{\phi(p, t)}{\log(1/p)} = \lim_{p \rightarrow 0} \frac{I_p(f_p)}{\log(1/p)} = \lim_{p \rightarrow 0} \frac{1}{2} \iint f_p(x, y) dx dy = \frac{(6t)^{2/3}}{2},$$

which proves the first assertion of the theorem. A second important consequence of (26), to be useful later, is that

$$(27) \quad \iint f_p(x, y)(1 - f_p(x, y)) dx dy \leq \frac{C}{\log(1/p)}.$$

Next, note that for any x, y ,

$$\begin{aligned} & \frac{1}{2} \iint (f_p(x, z)f_p(z', y) - f_p(x, z')f_p(z, y))^2 dz dz' \\ &= \iint f_p^2(x, z)f_p^2(z', y) dz dz' - h_p^2(x, y) \\ &= \int f_p^2(x, z) dz \int f_p^2(y, z) dz - h_p^2(x, y). \end{aligned}$$

It follows from this and (26) that

$$\begin{aligned} & \iiint (f_p(x, z)f_p(z', y) - f_p(x, z')f_p(z, y))^2 dz dz' dx dy \iint f_p^2(x, y) dx dy \\ & \leq \frac{C}{\log(1/p)}. \end{aligned}$$

The above inequality and the lower bound on $\iint f_p^2(x, y) dx dy$ from (26) give

$$\begin{aligned} (28) \quad & \iiint (f_p(x, z)f_p(z', y) - f_p(x, z')f_p(z, y))^2 dz dz' dx dy \\ & \leq \frac{C}{t^{2/3} \log(1/p)}. \end{aligned}$$

Let $M_p := \iint f_p(x, y) dx dy$. For each x , let $m_p(x) := M_p^{-1/2} \int f_p(x, y) dy$. An application of Jensen's inequality to (28) gives

$$\iint (f_p(x, z)M_p - M_p m_p(x)m_p(z))^2 dz dx \leq \frac{C}{t^{2/3} \log(1/p)}.$$

By (26), $M_p \geq (6t)^{2/3}$. Thus, m_p is bounded by $(6t)^{-1/3}$, and

$$(29) \quad \iint (f_p(x, z) - m_p(x)m_p(z))^2 dz dx \leq \frac{C}{t^2 \log(1/p)}.$$

For each p , let $n_p : [0, 1] \rightarrow [0, (6t)^{-1/3}]$ be a step function (i.e. a function that is constant on intervals) such that

$$(30) \quad \int (m_p(x) - n_p(x))^2 dx \leq p.$$

(Such functions exist because we can approximate m_p by a continuous function to any degree of accuracy by Lusin's Theorem [26], and then approximate the continuous function by a step function.)

Let σ_p be a measure preserving bijection of $[0, 1]$ such that $n_p(\sigma_p x)$ is a non-increasing function. Such a bijection is easy to construct because n_p is a step function. Let $\ell_p(x) := n_p(\sigma_p x)$ and $g_p(x, y) := f_p(\sigma_p x, \sigma_p y)$. By the monotonicity and uniform boundedness of ℓ_p there exists a sequence $\{p_i\}_{i \geq 1}$ decreasing to zero such that ℓ_{p_i} converges in L^2 to a limit function ℓ . Therefore by (29) and (30), $g_{p_i} \rightarrow g$ in L^2 , where

$$g(x, y) := \ell(x)\ell(y).$$

By this and (27), g is a $\{0, 1\}$ -valued function. It is not difficult to see from this and the non-negativity of ℓ that ℓ must also be $\{0, 1\}$ -valued. (If $\ell(x) \notin \{0, 1\}$ on some set of positive measure, then there may be a set A of positive measure where $\ell(x) \in (0, 1)$, or there may be a set A of positive measure where $\ell(x) \in (1, \infty)$. In either case, $g(x, y) \notin \{0, 1\}$ on $A \times A$.) Since ℓ is monotone decreasing, it follows that ℓ must be the indicator of an interval of the form $[0, b]$ for some $b \in [0, 1]$. Lastly, (26) implies that $\iint g(x, y) dx dy = (6t)^{2/3}$, and therefore $b = (6t)^{1/3}$. Consequently, $g = \chi_t$.

The above argument establishes that for any collection $\{f_p\}_{p>0}$ such that $f_p \in F_{p,t}^*$ for each p , there is a sequence $\{p_i\}_{i \geq 1}$ decreasing to zero such that $f_{p_i} \rightarrow \chi_t$ in the cut metric. The same argument can be extended to show that for any sequence $\{f_{p_i}\}$ such that $p_i \rightarrow 0$ and $f_{p_i} \in F_{p_i,t}^*$ for each i , there is a subsequence converging to χ_t in the cut metric. This proves the second assertion of the theorem. The last claim of the theorem follows from this and Theorem 4.1. \square

4.8. Open questions. There are many questions that remain unresolved, even in the simple example of upper tails for triangle counts that has been analyzed in this section. For instance, what is the set of optimal solutions of the variational problem (15) in the broken replica symmetry phase (i.e. where the optimizer is not a constant)? Is the solution unique in the quotient space $\widehat{\mathcal{W}}$, or can there exist multiple solutions? Is it possible to explicitly compute a nontrivial solution of (15) for at least some value of (p, t) ? Is it possible to even numerically evaluate or approximate a solution using a computer? Does Theorem 4.3 characterize the full replica symmetric phase? If not, what is the exact phase transition boundary? What happens in the sparse case where p and t are both allowed to tend to zero? At the time of writing this paper, we do not know how to answer any of these questions.

Acknowledgments. The authors thank Amir Dembo for his crucial role in motivating this research and suggesting the formula for the rate function in the case of subgraph counts, Joel Spencer for hinting that Szemerédi's lemma may be useful, and Persi Diaconis for pointing us to the work of Lovász and coauthors (which helped in cutting down the size of the paper by a half). The authors also thank the referee for a very careful report.

REFERENCES

- [1] ALDOUS, D. (1981). Representations for partially exchangeable arrays of random variables. *J. Multivariate Anal.* **11** no. 4, 581–598.
- [2] BOLLOBÁS, B. (2001). *Random graphs*. Second edition. Cambridge University Press, Cambridge.
- [3] BOLLOBÁS, B. and RIORDAN, O. (2009). Metrics for sparse graphs. *Surveys in Combinatorics, LMS Lecture Notes Series* **365**, CUP 2009, 211–287.
- [4] BOLTHAUSEN, E., COMETS, F. and DEMBO, A. (2003). Large deviations for random matrices and random graphs. *Private communication*.

- [5] BORGS, C., CHAYES, J., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2006). Counting graph homomorphisms. *Topics in discrete mathematics*, 315–371, Algorithms Combin., **26**, Springer, Berlin.
- [6] BORGS, C., CHAYES, J., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2008). Convergent sequences of dense graphs I. Subgraph frequencies, metric properties and testing. *Adv. Math.* **219** no. 6, 1801–1851.
- [7] BORGS, C., CHAYES, J., LOVÁSZ, L., SÓS, V. T. and VESZTERGOMBI, K. (2007). Convergent sequences of dense graphs II. Multiway cuts and statistical physics. *Preprint*. Available at <http://research.microsoft.com/en-us/um/people/borgs/papers/conright.pdf>
- [8] CHATTERJEE, S. (2010). The missing log in large deviations for subgraph counts. *Preprint*. Available at <http://arxiv.org/abs/1003.3498>
- [9] CHATTERJEE, S. and DEY, P. S. (2009). Applications of Stein’s method for concentration inequalities. To appear in *Ann. Probab.* Available at <http://arxiv.org/abs/0906.1034>
- [10] CHATTERJEE, S., DIACONIS, P. and SLY, A. (2010). Random graphs with a given degree sequence. *Preprint*. Available at <http://arxiv.org/abs/1005.1136>
- [11] DEMBO A. and ZEITOUNI, O. (1998). *Large Deviations Techniques and Applications*, Second Edition. Springer.
- [12] DIACONIS, P. and JANSON, S. (2008). Graph limits and exchangeable random graphs. *Rend. Mat. Appl. (7)* **28** no. 1, 33–61.
- [13] DIESTEL, R. (2000). *Graph theory*. Second edition. Springer-Verlag, New York, 2000.
- [14] ERDŐS, P. and RÉNYI, A. (1960). On the evolution of random graphs. *Publ. Math. Inst. Hungar. Acad. Sci.* **5** 17–61.
- [15] FREEDMAN, M., LOVÁSZ, L. and SCHRIJVER, A. (2007). Reflection positivity, rank connectivity, and homomorphism of graphs. *J. Amer. Math. Soc.* **20** 37–51.
- [16] FRIEZE, A. and KANNAN, R. (1999). Quick approximation to matrices and applications. *Combinatorica* **19** 175–220.
- [17] HOOVER, D. N. (1982). Row-column exchangeability and a generalized model for probability. *Exchangeability in probability and statistics (Rome, 1981)*, 281–291, North-Holland, Amsterdam-New York.
- [18] JANSON, S., ŁUCZAK, T. and RUCIŃSKI, A. (2000). *Random graphs*. Wiley, New York.
- [19] LOVÁSZ, L. (2006). The rank of connection matrices and the dimension of graph algebras. *Eur. J. Comb.* **27** 962–970.
- [20] LOVÁSZ, L. (2007). Connection matrices. In: *Combinatorics, Complexity, and Chance: A Tribute to Dominic Welsh* (ed. G. Grimmett and C. McDiarmid), 179–190. Oxford University Press, Oxford.
- [21] LOVÁSZ, L. and SÓS, V. T. (2008). Generalized quasirandom graphs. *J. Comb. Theory B.* **98** no. 1, 146–163.
- [22] LOVÁSZ, L. and SZEGEDY, B. (2006). Limits of dense graph sequences. *J. Combin. Theory Ser. B* **96** no. 6, 933–957.
- [23] LOVÁSZ, L. and SZEGEDY, B. (2007). Szemerédi’s lemma for the analyst. *Geom. Funct. Anal.* **17** no. 1, 252–270.
- [24] LOVSZ, L. and SZEGEDY, B. (2007). Testing properties of graphs and functions. To appear in *Israel J. Math.*
- [25] LOVÁSZ, L. and SZEGEDY, B. (2009). Contractors and connectors in graph algebras. *J. Graph Theory*, **60** no. 1, 11–30.
- [26] RUDIN, W. (1987). *Real and complex analysis*. Third edition. McGraw-Hill Book Co., New York.
- [27] SANOV, I. N. (1961). On the probability of large deviations of random variables. *Selected Translations in Mathematical Statistics and Probability* **1** 213–224. (Russian original: *Mat. Sb.* 1957.)

- [28] SZEMERÉDI, E. (1978). Regular partitions of graphs. *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, pp. 399–401, Colloq. Internat. CNRS, 260, *CNRS, Paris*.

COURANT INSTITUTE OF MATHEMATICAL SCIENCES, NEW YORK UNIVERSITY, 251
MERCER STREET, NEW YORK, NY 10012